

sEM Align

Copyright © 1999-2001 John C. Fiala
All rights reserved.

Table of Contents

<u>Purpose</u>	3
<u>System Requirements</u>	4
<u>Overview</u>	5
<u>File Organization</u>	6
Making Alignments	
<u>Section Movements</u>	7
<u>Flickering between Sections</u>	8
<u>Blending Sections</u>	9
<u>Aligning a Series</u>	10
<u>Aligning by Point Correspondences</u>	11
Menu Operations	
Series Menu	
<u>Open</u>	12
<u>Close</u>	13
Transforms	
<u>Apply To Images</u>	14
<u>Average</u>	15
<u>Adjust by New Transform</u>	16
<u>Adjust by View Transform</u>	17
<u>Absolute->Incremental</u>	18
<u>Incremental->Absolute</u>	19
<u>Restore</u>	20
Preferences	
<u>Use Working Images</u>	21
<u>Clear Working Images</u>	22
<u>Exit</u>	23
Section Menu	
<u>Go to</u>	24
<u>Restore from Disk</u>	25
<u>Restore Contrast</u>	26
<u>Import File</u>	27
<u>Mode: Antialias</u>	28
<u>Mode: Blend</u>	29
Transform Menu	
<u>Align Points</u>	30
<u>Clear</u>	31
<u>Restore from Disk</u>	32

<u>Restore from Backup</u>	33
<u>Auto Save</u>	34
<u>Save</u>	35
<u>Paste From</u>	36
<u>Type In</u>	37
View Menu	
<u>Adjustment Mode</u>	38
<u>Clear Offsets</u>	39
<u>Reorient</u>	40
<u>Fit Window to View</u>	41
<u>Use Window as View</u>	42
<u>Set Scaling</u>	43
<u>Scale Image to Fit</u>	44
<u>Copy to Clipboard</u>	45
<u>Background Color</u>	46
<u>Help Menu</u>	47

Purpose

sEM Align is a Windows program for aligning grayscale images. Align accepts only 8-bit per pixel Windows Bitmap (.bmp) or grayscale JPEG (.jpg) images as input and produces only 8-bit per pixel Windows Bitmap or JPEG grayscale output. The output is suitable for use with the program IGL Trace for making 3D measurements and reconstructions.

System Requirements

Minimum requirements:

100 MHz Pentium PC
8 Mb of free memory
800x600 display with at least 1024 colors

The speed of alignments is dependent on the speed of the computer and size of the on-screen view, so low system performance can be at least partially compensated by reducing the view size. During normal operation, the program keeps two scaled images in memory, and for each maintains additional an additional image buffer for the on-screen view. The program requires RAM memory sufficient to hold these images. For images scaled to fit the screen this is usually no more than 6 Mb. When Scaling is not used, memory requirements increase to that needed to hold the actual images in memory.

Recommended requirements:

200 MHz Pentium Pro or better
at least 64 Mb of memory
a fast 1024x768 display with 2Mb video memory

Overview

sEM Align facilitates the alignment of grayscale images by allowing on-screen adjustments of image position, orientation, scale, and distortions up to 2nd order. Alignments are made relative to other images in a sequence of images called a *series*. Each image in a series is referred to as a *section*, in reference to material obtained by serial sectioning tissue for electron microscopy, (sEM=serial electron microscopy). Thus, each section is an image stored in a separate file. The position of the section in the series is indicated by a numerical extension.

sEM Align is designed to allow on-screen alignment of larger images than is possible through the use of IGL Align and special purpose OpenGL hardware. sEM Align is implemented completely in software which runs on any general purpose PC. This is accomplished by displaying and moving scaled versions of the original images on screen. Speed of movement and display is proportional to the size of the on-screen view.

Section alignments can be made in several ways. Images can be incrementally moved using keyboard input. Alignments can be computed from a set of feature correspondences specified by the user. The resulting section *transforms* are stored in separate files. Thus, original image data is not modified by sEM Align. Additionally, the program allows the user to make adjustments to the entire series at once or to make adjustments to a separate view transform which gets applied to every image.

sEM Align can display two sections blended together. This facilitates putting features of two sections in correspondence since features in both sections can be seen while moving one of them. Another useful feature for validating alignments is flickering between two sections. Aligned features do not generate any apparent motion during flickering, while misaligned features generate apparent motion in the direction of the misalignment.

For increased performance, sEM Align normally does not antialias the display of transformed sections. It is not critical to do this during alignments since no changes are being made to the original data anyway. Each time a section is displayed, the transformed section is recomputed from the image data. However, antialiasing can be performed when desired.

File Organization

For each series, files are organized according to a particular naming convention based around a user-specified base name. The base name can be an arbitrary Windows 95-compatible name string up to 255 characters long. For example, suppose a base name of “Rat56-series35-Mar24-99” was selected for series of 100 sections. The file organization for this series would be as follows:

Section Images:

_Rat56-series35-Mar24-99.1
_Rat56-series35-Mar24-99.2
...
_Rat56-series35-Mar24-99.100

Section Transforms:

#Rat56-series35-Mar24-99.1
#Rat56-series35-Mar24-99.2
...
#Rat56-series35-Mar24-99.100

Old Section Transforms:

~Rat56-series35-Mar24-99.1
~Rat56-series35-Mar24-99.2
...
~Rat56-series35-Mar24-99.100

Working Section Images:

\$Rat56-series35-Mar24-99.1
\$Rat56-series35-Mar24-99.2
...
\$Rat56-series35-Mar24-99.100

The section images, once created and named according to this convention, can be in any directory which is readable, including CD-ROM or other read only media. This directory is called the *Images Directory*. The other files should reside together in a high-speed, writable directory such as on a hard disk drive. This is called the *Working Directory*. The user is only involved in creating the original section images, the other files are created and maintained by the program and do not need to be directly manipulated by the user. The section images may be placed in the *Working Directory*. All files must end with a purely numerical extension indicating the section number of the file.

Section Movements

A section can be aligned by incrementally adjusting its position, orientation, scale, slant and deformation. These movements are easily made using the keyboard. The following keys move the current section one increment. The size of the increment is multiplied ten-fold by simultaneously holding the control (Ctrl) key down. The size of the increment can be adjust in the Transform:Type In dialog.

Up Arrow	Move section upward.
Down Arrow	Move section down.
Left Arrow	Move section to left.
Right Arrow	Move section to right.
c	Rotate section clockwise.
z	Rotate section counterclockwise.
F1	Slant horizontally in counterclockwise direction.
F2	Slant vertically in counterclockwise direction.
F3	Slant vertically in clockwise direction.
F4	Slant horizontally in clockwise direction.
F5	Deform to reduce top of section horizontally.
F6	Deform to reduce right side of section vertically.
F7	Deform to expand right side of section vertically.
F8	Deform to expand top of section horizontally.
F9	Scale to shrink section horizontally.
F10	Scale to shrink section vertically.
F11	Scale to enlarge section vertically.
F12	Scale to enlarge section horizontally.

By default the center of rotation of a section is its lower left corner. Clicking the right mouse button on the current section sets a new center for subsequent rotations, slants, deformations and scalings of that section.

Flickering between Sections

The current section can be changed by using the Section:Go to command, or by using the Page Up and Page Down keys. When a new section is chosen by one of these methods the current section becomes the previous section and the new section becomes the current section.

A useful means for checking the alignment between two sections is to rapidly switch (flicker) the display of sections on the screen. Aligned features do not generate any apparent motion during flickering, while misaligned features generate apparent motion in the direction of the misalignment. Flickering can be easily done by holding down the '/' key. This switches the displayed image between the current and previous sections.

After observing the apparent motion during flickering, refine the alignment using the section movement commands from the keyboard. Be careful that the adjustments are only made to the desired section, since the program will allow either image to be moved.

Blending Sections

Although flickering is useful for checking alignments, it is often easier to perform an alignment in blend mode. When in blend mode the current section is displayed simultaneously with the previous section. Movements of the current section can be made using the keyboard commands until the desired alignment is realized. Switching between blend and flicker mode can be helpful to monitor the progress of the alignment. Be careful to reenter blend mode from the correct section, such that subsequent movements will be applied to that section.

Aligning a Series

To align a series of sections, first digitize and place the section images into the proper file organization. Start sEM Align and adjust the window size to the desired working dimensions. Open the series and Goto a section from which to begin the alignment. Save the transform. This will create a transform file for the section which contains the identity transform. Page Up (Down) to the adjacent section in the series. Compare the section to the previous section using flickering or blending. Align the current section to the previous section using one of the methods described below.

It may be helpful to save intermediate alignment attempts before making additional refinements. If the refinements don't work out, the intermediate alignment can be easily restored from disk rather than starting over. When alignment is complete, save the transform. Proceed to the next section and align it to the one just completed. Continue until the entire series has been aligned.

Absolute Alignment:

Alignment creates a transform for each section. This transform can be used in one of two ways: as a relative transformation between adjacent sections or as an absolute positioning of the section into the aligned sequence. In this latter (*absolute*) mode each section is aligned to the previously aligned adjacent section such that the all sections are eventually aligned together.

Incremental Alignment:

Sometimes aligning in absolute mode introduces a systematic shift or scaling of images. For example suppose section 1 is not scaled, section 2 is scaled by 102%, section 3 is scaled by 102%, section 4 is scaled by 102%, and so on. In this case section 18 would turn out 40% bigger than section 1. This effect can be avoided by using an *incremental* alignment. In incremental mode each section is aligned to the untransformed adjacent section. After making an alignment for a section (call it section 1), save the transform, page to the next section (2). Flip back ('/') to section 1 and clear the transform. Now flip to section 2 and align it to the untransformed section 1. Save the resulting transform for section 2. Then page to the next section and continue. However don't save the cleared section 1 transform as this isn't part of the alignment!

Aligning by Point Correspondences

If there are correspondences between particular features in the images to be aligned, the program can compute the transform necessary to bring them into alignment. The user only needs to specify which points correspond.

Each section, when it is loaded, has the capacity of remembering an arbitrary number of feature points on the image. These points are specified by the user by clicking the left mouse button when the tip of the cursor is at the desired point. Points can be deleted by clicking on them a second time. Points for a section are stored in the order they are entered. For purposes of alignment, the first point entered on the current section corresponds to the first point entered on the previous section, the second point to the second, and so on.

Point correspondences entered by the user can be used by the program to automatically align the sections. These computed alignments can be in either absolute or incremental mode, as selected on the Transform:Align Points menu. Computed alignments overwrite any other transform parameters previously input.

To align by point correspondences, Page to the next (unaligned) section. Select a feature point in this section. Hit the '/' key to switch to the previous section. Select a corresponding feature point in the previous section. Switch back to the unaligned section using the '/' key. Enter a second feature point. Switch to previous and enter the corresponding point. When sufficient correspondences have been made, switch back to the unaligned section and select one of the alignment options from the Transform:Align Points menu, e.g. Linear. The current section will be redisplayed with the computed alignment. Hold down the '/' key to verify that the selected points are aligned. There should be little or no apparent motion between the correspondences.

If the alignment is inadequate, try selecting a more complex alignment option that includes higher-order terms. If new correspondences need to be added, click on additional feature points and reselect the alignment operation from the Transform:Align Points menu. If the alignment is completely wrong try clearing the alignment points and entering new ones.

Open

This operation allows the user to open an existing series or to create a new series. The user is presented with a dialog box that asks for the base name of the series to open. The user can change the drive and directory to the appropriate location. When a directory contains identifiable components of a series, e.g. the section image files or the transform files, the base name of the series is shown in the dialog box. The series can be selected by double clicking the series name with the left mouse button.

Alternatively, the user may type a new base name to create a new series. Valid base names can be up to 255 characters long, using any characters valid in a Windows 95-compatible file name.

When a series is opened, the directory shown in the dialog box is set to be both the *Images Directory* and the *Working Directory*. If, later, this directory is found to be unwritable, the user will be asked to enter a new *Working Directory*. If the directory is writable the user may select a different *Images Directory* in the Section:Go to dialog.

Close

The menu item closes the currently open series. It is not necessary to close a series before opening another one, since the current series will be automatically closed by opening another series.

Apply To Images

This item of the Transforms submenu generates a new series of images from a subset of the section images. The user is prompted for a set of parameters which define how the new series is to be generated. Then each section image is read, transformed using both the view and section transforms with antialiasing. Then cropped and stored to disk as a new file. The following describes how each parameter is used:

Section Range:

First Section: The first section to process.

Last Section: The last section to process.

Transformation Target Size:

Scale: The amount of scaling to apply when reading the original image data into memory. This amount is fixed at 1.0 to minimize aliasing of original data. Any global size changes can be realized through the scale of the view transform.

Width: The width of the transformed image (into which the scaled image will be mapped.)

Height: The height of the transformed image.

Subregion of Target to Extract:

Lower Left: X,Y: The lower left corner of a rectangle in the transformed image from which the final image will be cropped.

Width: Width of rectangle to be extracted.

Height: Height of rectangle to be extracted.

Destination:

Path: The directory in which the output files will be created. To browse to a new location, click on the Path button.

Series: The base name of the new series to be generated.

To make sure that all sections are completely contained in the output use this procedure:

In View Mode, find the section with the maximum downward translation. Adjust the view until its lower edge is visible. Similarly, adjust the view so that no sections are outside the view to the left. This ensures that no part of the series lies below the lower left corner (0,0). Now select the Apply To Images menu item. The Apply To Images dialog will automatically be filled with default values for target width and height such that every section is completely contained in the resulting target image. Using these default values makes all sections completely visible in the output, since they are now contained in the region from (0,0) to the target width and height. Test the parameters on a single image and verify that the output is as desired before running on the entire series.

This operation can be very slow, particularly if the section images are large. To abort the operation, hold down the Control key and press End. The operation will be stopped after a delay but before the next section in the series is processed.

Average

This item of the Transforms submenu allows the user to average the linear part of the transforms in a particular section range. The user is prompted for the first and last section in the range to be averaged. The linear part of the transforms of sections in this range are summed and divided by the number of sections. The inverse of the resulting average linear transform is then be applied to each section in the range, resulting in a new set of transforms of reduced magnitude. The modified sections retain their original alignment, but have an altered alignment with any sections outside of the selected range.

Before modifying the transform of a section, this operation first saves the transform in a backup file. This allows the original transforms to be recovered using the Transforms:Restore command.

Adjust By New Transform

This item of the Transforms submenu allows a set of transforms to be modified by a user specified amount. The user is first prompted for the range of sections to be modified. Then the user is asked to type in a particular transform using the standard Transform Type In dialog. This dialog includes the fields for the *Amount* of individual transform components and fields for *Keystroke* increments. Changing the values in the *Keystroke* column will have no effect on the adjustment operation, however these changes will be used in subsequent image movements made from the keyboard, such as shifts with the arrow keys, for example.

Since the same transformation is applied to all sections in the specified range, this command does not alter the alignment within that range. Nonlinear terms in the adjustment transform are ignored.

Adjust By View Transform

This item of the Transforms submenu allows a set of transforms to be modified by the view transform. The user is first prompted for the range of sections to be modified. Then the view transform is applied to each section and the result saved as the new section transform. Since the same transformation is applied to all sections in the specified range, this command does not alter the alignment within that range. This command provides a means for saving a desired view, by incorporating it into the individual section transforms.

Incremental -> Absolute

Use this operation to convert incremental transforms to absolute transforms. The user is first prompted for the range of sections to be modified and the direction of the incremental alignments. Then each section transform is converted to an absolute alignment transform by accumulating the total transform through the series. The computed absolute transforms are saved to disk and the previous incremental transforms are moved to the transform backup files. The incremental transforms can thus be recovered by using the Restore command, but only until the next update of the transform files.

Note that since this operation uses numerical methods, it is not recommended to switch back and forth between absolute and incremental alignment modes repeatedly. Repeated use of this operation could lead to an accumulation of numerical errors that might introduce misalignments.

Absolute -> Incremental

Use this operation to convert absolute transforms to incremental. The user is first prompted for the range of sections to be modified and the desired direction of the incremental alignments. Then each section transform is converted to an incremental alignment transform by removing the adjacent transform's contribution to the section transform. The computed incremental transforms are saved to disk and the previous absolute transforms are moved to the transform backup files. The absolute transforms can thus be recovered by using the Restore command, but only until the next update of the transform files.

Note that since this operation uses numerical methods, it is not recommended to switch back and forth between absolute and incremental alignment modes repeatedly. Repeated use of this operation could lead to an accumulation of numerical errors that might introduce misalignments.

Restore

This item of the Transforms submenu restores the section transforms from their last backup. A backup is written whenever a section transform is saved, e.g. after an averaging operation or after an individual Save command.

Use Working Images

Working images are .bmp image files used to store the scaled version of the original image data. Since the original image data file can be very large (>10Mb) while the scaled version used for alignment is relatively small (<1Mb), performance can be improved by saving a working copy of the scaled image rather than reloading the original image data each time a section is accessed. This is only done when the user has enabled working images by selecting this item of the Preferences submenu.

When accessing a section, the *Working Directory* is searched for a corresponding working image. If it is found and matches the current scaling, it is loaded instead of reading and scaling the section image. If a proper working image is not found, the image data is read from the section image and a new working image is written. In order to recover the original image size, working images use the `biXPelsPerMeter` and `biYPelsPerMeter` fields of the .bmp info header to store the original image width and height.

Clear Working Images

This item of the Preferences submenu allows the user to remove from the *Working Directory* any working images for a series.

Exit

Terminate the sEM Align application. It is not necessary to close an open series before exiting the program.

Go to

This command lets the user select a section from the series, after which the current section becomes the previous section and the section selected by the user becomes the current section. The section image and transform are loaded from files and the transformed image is displayed in the window.

The user may also browse to a new directory which contains the section images for the series. This new directory becomes the *Images Directory*. The *Working Directory* is not modified by this command.

Restore from Disk

This command reloads the current section from disk, resampling the image at the working scale factor and reloading the transform from the file. This command is the same as Section: Go to the current section.

Restore Contrast

This item restores the current section's contrast and brightness settings to their original values. Brightness and contrast can be manipulated using the keyboard. The '[' key increases contrast while the '[' key decreases contrast. The '=' key increases brightness while the '-' key decreases brightness.

Brightness and contrast settings are saved in the transform file for each section. Thus, they do not modify the original image data.

Import File

This command copies an image file to a particular section in a series. The new file is named to indicate the series base name and section number as described in the file organization.

Either 8-bit Windows Bitmaps or grayscale JFIF (JPEG) image files can be used as section images. To convert other formats, use the Convert program available at <http://synapses.bu.edu/tools/>

Mode: Antialias

This submenu item toggles the state of antialiasing. When the item is checked antialiasing is enabled. All operations subsequently use the slower method of antialiasing to generate the on-screen transformed section images. This item can also be toggled by using Ctrl-A on the keyboard.

Mode: Blend

This submenu item toggles the state of blending. When the item is checked blending is enabled. The current image and the previous image are displayed on-screen blended together. This item can also be toggled using Ctrl-B on the keyboard.

Align Points

This submenu has functions for automatically generating alignments based on corresponding points. Alignment points are selected within a section by placing the tip of the cursor on the desired point in the section and clicking the left mouse button. The first such point entered on the current section corresponds to the first point entered on the previous section. The second point to the second, and so on. When sufficient correspondences have been made, selection of the following functions will generate a new transform for the current section.

Rigid:

Compute a rigid-body transform that minimizes the error between corresponding points. Rigid transforms include translation and rotation only. Rotational adjustments computed by Rigid are limited to ± 90 degrees. At least 2 pairs of corresponding points are required to perform this computation.

Linear:

Compute a linear transform that minimizes the error between corresponding points. Linear transforms include translation, rotation, scale, and slant. At least 3 pairs of corresponding points are required to perform this computation.

Deformal:

Compute a transform with linear and xy terms that minimizes the error between corresponding points. The xy terms introduce deformation, i.e. scaling one side of the image more than the other side. At least 4 pairs of corresponding points are required to perform this computation.

Quadratic:

Compute a full quadratic mapping that minimizes the error between corresponding points. The terms include linear, xy, as well as x^2 and y^2 terms. At least 6 pairs of corresponding points are required to perform this computation.

Mode:

Select the alignment mode for point correspondence computations. In **absolute** mode a section is aligned to the transformed previous section. In **incremental** alignment mode a section is aligned to the untransformed previous section.

Point Color:

Select the color which new points will be given. Once a point is created, its color is fixed.

Clear Points:

This item clears all the points from both the previous and current section.

Clear

This item sets the current transform to the identity (i.e. no) transform.

Restore from Disk

This submenu item reloads the current transform from the file. Any changes in the current transform will be lost.

Restore from Backup

This submenu item discards the current transform and the current transform file. It replaces the current transform file with the most recent backup version of the transform file, and loads this new transform file into memory.

Auto Save

This item toggles Auto Save mode. When in Auto Save mode, changes to transforms are automatically saved to disk without prompting the user. When not in Auto Save mode, the user is asked whether a changed transform should be saved before deleting it from memory.

Save

This command saves the current transform to a file. The old transform file is copied to a backup file, from which it can be restored if necessary.

Paste From

This menu item allows the user to copy a transform from another section to the current section. The user is prompted for the section number. This transform replaces the current transform in memory.

Type In

This menu item invokes the Transform Type In dialog. In this dialog the user can modify the current transform as well as the increments used for on-screen movements, such as shifts with the arrow keys. The current transform is displayed in the fields of the *Amount* column and in the center-of-rotation fields. Changing any of these values changes the current transform accordingly. The size of increments made by keyboard movements are displayed in the *Keystroke* column. Changing the values in the *Keystroke* column will have no effect on the current transform, however these changes will be used in subsequent image movements made from the keyboard.

Adjustment Mode

The transformation of each section for display on the screen involves two transforms, the transform of the individual section and a view transform which is applied to every section in the series. This menu item toggles which transform receives keyboard adjustments: the current transform or the view transform. When in view mode, the transform menu is not available and all keyboard movements apply to the view transform. This menu command can also be invoked by hitting the space bar on the keyboard.

Clear Offsets

This menu item clears the view transform.

Reorient

This command computes a new view transform based on the correct section transform. The new view transform is the inverse of the linear part of the current section transform. This results in the current section appearing in its original alignment, except for the effects of nonlinear terms.

Fit Window to View

This item changes the window size to hold exactly the current view. Note that the current view is the target rectangle into which the current image is transformed prior to being displayed on the screen. It can be bigger or smaller than the current image and bigger or smaller than the current window. Mismatches between the window and view size can be caused by resizing the window.

Use Window as View

This menu item causes the window size to be used as the size of the view. The view being the target into which the transformed image is mapped. When the window is resized by dragging a corner or side of the window frame, the view size is not automatically updated to match. Thus, parts of the transformed image may lie outside the view and not be visible, even though the window is big enough to reveal them. To remedy this problem, select this menu item after each window resizing operation.

Set Scaling

When a section image is loaded from a file, it can be scaled into a smaller image to fit on the screen and save memory. Note that this scaling occurs prior to the application of the both the section transform and the view transform. Scaling is different from transform scaling. Since every image in a series is displayed using the same scale factor, it does not alter the alignment. Global scaling can also be achieved through the adjusting the scale of the view transform. The advantage of using the scale factor during initial image loading is that less computer memory is used for the section image.

A scale factor of 1.0 is equivalent to no scaling, the image in memory is the same size as in the original file. A scale factor of 2 means that the dimensions (width and height) of the memory object are half the size of the original image. The scale factor is user selectable by two submenu items:

None:

Sets the scale factor to 1.0.

Custom:

The user is prompted for a new scale factor. This becomes the scale factor for the display of images in the current series.

Scale Image to Fit

This command computes a new scale factor for the series and adjusts the view transform such that the entire transformed section image fits in the window. The requisite scaling change is made by adjusting the scale factor, not by adjusting the scale of the view transform. The view transform adjustment is to translate the image into the view window.

Copy to Clipboard

This command copies the current view (whether of the current section or of two blended sections) to the clipboard. The clipboard image can then be pasted into other applications.

Background Color

This menu item allows the user to enter a new grayscale value for the window background color. By default the background is black (a grayscale value of zero). Any grayscale value from 0 to 255 may be entered.

Help Menu

Special Keys:

This menu item displays a Windows help file containing the keyboard and mouse commands used for special alignment functions.

Help Align:

This command invokes this help file.

Image Info:

This menu item displays the size of the original image for the current section, as well as the size of the image in memory and the corresponding scale factor. Also displayed are the size of the view image and the size of the window.

About:

This menu item displays version, authorship and copyright information.